



Europäisches Patentamt
European Patent Office
Office européen des brevets



Publication number:

0 684 719 A1

12

EUROPEAN PATENT APPLICATION

21 Application number: **94480047.3**

51 Int. Cl.⁶: **H04L 29/06**

22 Date of filing: **25.05.94**

43 Date of publication of application:
29.11.95 Bulletin 95/48

64 Designated Contracting States:
DE FR GB

71 Applicant: **International Business Machines Corporation**
Old Orchard Road
Armonk, N.Y. 10504 (US)

72 Inventor: **Calvignac, Jean**
187 Chemin les Vallières
F-06610 La Gaude (FR)
Inventor: **Giroir, Didier**
Résidence Aurore,
84 Chemin du Val Fleuri
F-06800 Cagnes Sur Mer (FR)
Inventor: **Galand, Claude**

56 Avenue des Tuilières
F-06800 Cagnes Sur Mer (FR)

Inventor: **Lebizay, Gerald**

113 Avenue des Poilus

F-06140 Vence (FR)

Inventor: **Mauduit, Daniel**

Jardin des IRIS A5,

1 Chemin de la Batterie Russe

F-06200 Nice (FR)

Inventor: **Spagnol, Victor**

12 bis Chemin des Collettes

F-06800 Cagnes Sur Mer (FR)

74 Representative: **Vekemans, André**
Compagnie IBM France
Département de Propriété Intellectuelle
F-06610 La Gaude (FR)

54 Method and apparatus for transmission of high priority traffic on low speed communication links.

57 In a packet switched communications system an incoming real-time packet is imbedded after the next block of data of the non-real-time packet being transmitted.

This object is accomplished by transmitting each

packet along with at least a 1-byte trailer which is used to indicate the packet type, whether the current block of non real time data is preempted or whether the current block of non real time data is resumed.

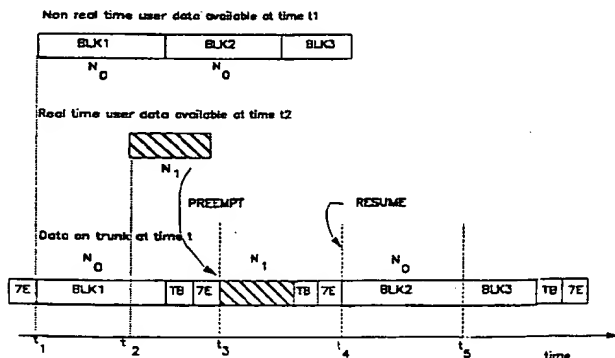


FIG. 2

Field of the Invention

This invention relates to telecommunications networks and more particularly to an improved preempt/resume protocol for low-speed communication links in a fast packet switching network.

Background of the Invention

In a digital transmission network, data from a large number of users are serially transmitted from one network node to another network node, up to their respective final destination.

Due to the evolution of networks towards more and more complex mixings of sub-networks with heterogeneous architectures, it is clear that there is a future requirement to support distributed computing applications across high speed backbones that may be carrying LAN traffic, voice, video and traffic among channel-attached hosts and work stations. Perhaps the fundamental challenge for high speed networking is to minimize the processing time within each node in the network.

Packet switching is now commonly used to accommodate the bursty, multiprocess communication found in distributed computing environments. To accomplish this, packets carrying bursty data traffic can be assigned a non-real-time priority, while packets carrying voice and video traffic can be assigned a higher, real-time priority. A node in a fast packet switching network contains buffers for holding packets waiting for transmission on its communication links. Packets waiting for transmission can be held in buffers managed differently, depending on the priority, assigned to the packets.

European patent application (582537) discloses a communication network having communication nodes which can adopt a number of different service policies in order to transmit packets from different priority buffers, for instance, priority with no preemption, preemption with retransmission, and priority with resume. When no preemption is used, the packet priority is only examined to determine from which buffer to select the next packet for transmission. If a high-priority packet is placed in the buffer while a low-priority packet is being transmitted, the high-priority packet must wait until the current transmission is completed. A preemption with retransmission service policy means that the node will abort the transmission of a low-priority packet upon the arrival of a high priority packet and transmit the high-priority packet. Once all high-priority packets have been transmitted, transmission of the preempted low-priority packet will be restarted from the beginning of the packet. A preemption with resume service policy is similar except the preempted low-priority packet is restarted from the point of interruption rather than the

beginning. The selection of the appropriate service policy is dependant on the characteristics of the communication link, the delay requirements of the high-priority packets, and the size of the low-priority packets.

The typical scheme used for transmitting packetized information over low speed communication links (T1 at 1.544 megabits per second) is based on the HDLC MAC-layer protocol, described, for example, in H. NUSSBAUMER: TELEINFORMATIQUE I, Presses Polytechniques Romandes, 1987, pages 301-313.

Both the priority with no preemption and the preemption with retransmission service policies can be implemented using the existing HDLC MAC-layer protocol. For preemption with resume service policy, a modified HDLC MAC-layer protocol is described wherein three types of flags are used to delimit packets for allowing high-priority packets to temporarily preempt low-priority packets. The HDLC starting, ending or idle flag is defined as the 8-bit sequence B'01111110'(X'7E'). The start preempt flag is defined as the 9-bit sequence B'011111110' and the end-preempt flag is defined as the 10-bit sequence B'0111111110'. All flags are on byte boundaries with respect to the packet data that they delineate.

This definition requires that the hardware is capable to scan the incoming bit stream, to recognize special non-standard flags in addition to the HDLC flags, and to run a protocol to verify a set of rules upon detection of these flags. Clearly, a special hardware is necessary for that purpose.

Objects of the Invention

It is therefore a principal object of the invention to provide a method and an apparatus for embedding high-priority traffic in low-priority traffic for serial transmission through low speed communication links, without delay incurred by having first to complete transmission of a low-priority traffic.

It is another object of the invention to preempt low-priority packet traffic with at least one high priority packet and later resume transmission of the preempted low-priority packet automatically, using an output adapter built with off-the-shelf scanners.

It is another object of the invention to allow a mode of operation that is compatible with the HDLC MAC-layer protocol.

Brief Summary of the Invention

According to the invention an incoming real-time packet is imbedded after the next block of data of the non-real-time packet being transmitted.

This object is accomplished by transmitting each packet along with at least a 1-byte trailer

which is used to indicate the packet type, whether the current block of non real time data is pre-empted or whether the current block of non real time data is resumed.

Brief Description of the Drawings

The invention will be described with respect to a preferred embodiment thereof, which is further shown in the drawings in which:

Figure 1 is a block diagram representing the output adapter of a node of a network within which the invention may be practiced.

Figure 2 represents the principle of the pre-empt/resume mechanism according to the invention.

Figure 3 represents the configuration of a trailer byte according to the invention.

Figure 4 shows a point to point T1 line connection.

Figure 5 is a model of the different transmit sequences.

Figure 6 to 10 are different transmit sequences.

Figure 11 is a Preempt/Resume Transmit flow diagram.

Figure 12 is a Preempt/Resume Receive flow diagram.

Figure 13 is a block diagram of the transmit portion of a trunk interface of a communication system to which the invention is applicable.

Figure 14 is a block diagram of the receive portion of a trunk interface of a communication system to which the invention is applicable.

Detailed Description of Preferred Embodiment of the Invention

A packet-switching network usually consists of switching nodes and communication links or trunks between these nodes. Each of those nodes is linked to at least one other node through one or more communication links. The switching nodes are data processing systems including transmit/receive adapters connected to the communication links. At each node, incoming data packets are selectively routed to one or more of the outgoing communication links terminated at another node. Such routing decisions are made in response to information in the header of the data packet.

In a packet switching network, packets are pieces of data, which are prefixed with headers containing control and routing information that identifies the originating and destination users. Each node examines each header and decides where to send the packet to move it closer to its destination.

A basic requirement of high speed networks is to selectively process data according to different classes of services, which are generally specified

in terms of probability of loss and maximum end-to-end delay. This class of service may be specified by some bits in the header, which are decoded at intermediate nodes to select the buffering policy.

Delay priorities are specified among three classes of traffic. At each output trunk adapter, packets from each class share a different logic buffer before being transmitted over the trunk. These classes are:

- Real-time traffic (voice, video),
- Non-real-time traffic (data),
- Non-reserved traffic (datagram).

Figure 1 represents the output adapter structure, featuring three buffers 10, 11, 12, a scheduler 20, and a scanner 30. Packets received from the switch are stored in one of the buffers, according to their class, and the scheduler implements a policy to forward these packets to the output trunk. The scanner implements the low-level of the DLC protocol (frame synchronization, 0 insert/delete, CRC). The scheduler implements the preempt/resume protocol which will be described later.

Real-time traffic is given priority over non-real-time traffic in order to reduce its delay. Both real-time and non-real-time traffics are given priority over non-reserved traffic in order to minimize the impact of this class on the band-width reservation mechanism. For the sake of convenience, both non-real-time traffic and non reserved traffic will be considered as low priority traffic. Within each class, the packets are served in sequence, that is in the same order they arrived. Depending on the trunk speed, the scheduling policy is either preemptive or nonpreemptive.

In the case of non-preemptive policy, the buffer with the lower priority class is served only if the buffer with the highest priority class is empty, and the service of the low priority packets is not interrupted even when a high priority packet arrives before the end of the service.

This policy is used on all links for which the service time of a maximum-length non-real-time packet is less than 1.5 ms. For example, a T3 link supporting 2 Kbyte maximum length packets would use a nonpreemptive policy (0.3 ms maximum service time).

In the case of a preemptive resume policy, the buffer with the lower priority class is served only if the buffer with the highest priority class is empty, and the service of the low priority packets is interrupted when a high priority packet arrives before the end of the service. The service of the low-priority packet is resumed after the high-priority packet has been served.

This policy is used on all links for which the service time of a maximum-length non-real-time packet exceeds 1.5 ms. For example, a T1 link supporting 2 Kbyte maximum length packets would

use a preemptive with resume policy (10 ms service time).

This low speed interface uses the standard HDLC protocol: the X'7E' flag indicates the starting/ending of a packet. This flag is a shared starting and ending flag, i.e. one flag between packets, that defines byte alignment and is sent during idle periods. Having the X'7E' define byte alignment means that all packets using the network basic packet format are an integer number of bytes in length.

The HDLC receiver looks for X'7E' to synchronize the starting/ending of a packet. To assure that this bit pattern does not appear in the packet, zero bit stuffing is used. The HDLC transmitter inserts an extra 0 bit after each occurrence of five 1's in the data to be transmitted. The HDLC receiver monitors the bit stream for a pattern of five 1's. If the sixth bit is a zero, then it is deleted. If the sixth bit is a one, the seventh bit is a zero, and it is byte aligned, then this is a X'7E' flag. If the X'7E' flag is not byte aligned, then it is an abort flag, a non-aligned X'7E'.

A node must always be able to send and receive link-initialization and link liveness messages. These messages are only sent using the basic packet format for low-speed bit level interfaces.

The decision of whether to use preempt/resume protocol and HDLC frame format are determined during link-initialization. If both nodes agree to support connection oriented LLC, the HDLC frame format is used. In the HDLC frame format, data is sent using Information frames and link state transition exchange is sent using Unnumbered Information frames.

In the system described in the prior art, the preempt/ resume protocol uses a couple of special flags:

'011111110': Starting preempt flag

'011111110': Ending preempt flag

Both flags are byte aligned.

Let's assume that a non-real-time packet is ready to be transmitted at time t_1 , and that no real-time packet is ready to be transmitted. Then, the transmission of the non-real-time packet can begin at time t_1 . Now assume that a real-time packet arrives at the output trunk at time t_2 before the end of the transmission of the non-real-time packet. The packets are sent from the queue to the trunk adapter on a byte basis, and a real-time packet can be inserted after any byte of a non-real-time packet, provided it is encapsulated between a starting preempt flag (SP) and an ending preempt flag (EP). During the reception of the packet, the receiver looks for the starting preempt flag. Upon detection of this flag, it checks the validity of the detection, and eventually starts receiving the real-

time packet that has been imbedded within the non-real-time packet. It also continues looking at the received bits, detects the ending preempt flag, then checks the validity and resumes the reception of the non-real-time packet.

The validity of the preemption is defined by the following set of rules:

- X'7E' defines byte alignment.
- Six 1's preceded by a 0 that is not byte-aligned is an invalid code.
- Nine 1's preceded by a 0 that is not byte-aligned is an invalid code.
- Receipt of an invalid code aborts the current packet and subsequent packets, until X'7E'.
- Must verify that the preempted packet is a non-real-time packet (C1 bit of control byte 1 in the network header must be 1).
- Must verify that the packets received during preemption are real-time packets (C1 bit of control byte 1 in the network header must be 0).
- Cannot preempt a non-real-time packet before first byte is transmitted, because first byte is used to determine if the packet is real-time or non-real-time.

The invention will now be described with reference to Figure 2. It is easy to observe that in any case, the real-time traffic can support a small delay before being scheduled, provided this delay is bounded by a small maximum delay T . For example, one can consider that for a voice connection a maximum 1.5 ms delay per node is affordable, when compared to the 100 ms maximum end-to-end delay that is usually agreed on. Under this assumption, the preemption granularity can be defined at a block level instead of a byte level. In other words, instead of imbedding an incoming real-time (RT) packet after the next byte of the non-real-time (NRT) packet being transmitted, one can schedule its transmission for the next block of data of the packet being transmitted.

On trunks that support preempt/resume protocol, each packet is transmitted along with at least a 1-byte trailer which is used to indicate whether the packet is preempted. Each preempt/resume transition is marked by the transmission of a X'7E' flag preceded by a 1-byte trailer.

NRT (and NR) packets are segmented into short blocks ($N_0 = 128$ bytes). The scheduler sends each block over the line after it has checked that the RT queue is empty. At T1 speed, the block length corresponds to a 0.7 microsecond transmission delay.

If a RT packet arrives after the transmission of a NR or NRT packet has started, the scheduler notices the presence of that packet at the next inter-block checking. It then inserts a 1-byte trailer which indicates that this packet is going to be

preempted by at least one RT packet, triggers the transmission of a X'7E' flag and inserts the RT packet.

After the RT packet has been transmitted, the scheduler inserts a 1-byte trailer which indicates that this packet was a RT packet, triggers the transmission of a X'7E' flag, and checks again the RT queue. If the RT queue is empty, the scheduler resumes the transmission of the NRT (or NR) packet. Else, it transmits the next RT packet and then inserts a new 1-byte trailer and a X'7E' flag.

Figure 2 shows the principle of the preempt/resume protocol. Let's assume that the transmission of a non-real-time packet has started at time $t = t_1$, and that a N 1 byte real-time packet arrives at time $t = t_2$. The real-time queue is checked before sending a new block of the non-real time packet, so in our example the preemption will be activated at time $t_3 = t_1 + T$. After the real-time packet has been transmitted, the transmission of the non-real-time packet is resumed by blocks of N 0 bytes.

The trailer byte actually includes

- Only three state control bits that are used by the preempt/resume protocol.
- One reserved bit and
- Four other bits used for protection.

Figure 3 shows a trailer byte wherein:

B0 to B3:	Trailer byte checksum. Equal to the inverted value of bits B4 to B7.
B4:	Reserved
B5:	PT (Packet Type) bit: 0 Real-Time 1 Non-real time / Non reserved
B6:	PB (Preempted Block) bit: NRT (B5 = 1) 0 Current block is not preempted 1 Current block is preempted RT (B5 = 0) 0 Not applicable 1 Always inserted RT packet
B7:	RB (Resumed Preempt) bit: NRT (B5 = 1) 0 Current block is not resumed 1 Current block is resumed RT (B5 = 0) 0 Not last RT packet; next = RT 1 Last RT packet; next = NRT

Instead of using 4 parity bits, one could use an error correcting code which would correct all single errors in a trailer byte.

The different Preempt/Resume sending sequences will now be described with reference to Figures 4-10.

The data sequences described hereafter are those generated at the transmit side (sender trunk) then received at the other side (receiver trunk) thru a point to point T1 line connection, as shown in

Figure 4.

In the following text and figures

- NRT stands for non-real-time and non-reserved types of data packets.
- TB stands for Trailer Byte.
- F stands for X'7E' delimiter flag.

Data sequence models are represented according to the model of Figure 5

All possible sending sequence types generated by the sender are:

- Single NRT packet not preempted.
- Single RT packet not inserted.
- Preempted NRT packet, single preemption with one RT packet insertion.
- Preempted NRT packet, single preemption with several RT packets insertion.
- Preempted NRT packet, multiple preemption.

The sequence of Figure 6 represents a Single NRT packet, where T100 TB means: NRT complete packet, i.e. not pre emptied, not resumed.

The sequence of Figure 7 represents a Single RT packet where T000 TB means: RT PACKET, not inserted.

The sequence of Figure 8 represents a NRT packet preempted ONCE with one inserted RT packet, where T110 TB means: NRT, preempted, FIRST BLOCK of packet, T011 TB means: RT PACKET, inserted, followed by a NRT BLOCK, T101 TB means: NRT, resumed, LAST BLOCK of packet.

The sequence of Figure 9 represents a NRT packet preempted ONCE with several inserted RT packets, where T110 TB means NRT, preempted, FIRST BLOCK of packet, T010 TB means RT PACKET, inserted, followed by another RT PACKET, T011 TB means RT PACKET, inserted, followed by a NRT BLOCK, T101 TB means NRT, resumed, LAST BLOCK of packet.

The sequence of Figure 10 represents a NRT packet preempted SEVERAL times (SEVERAL inserted RT packets), where T110 TB means: NRT, preempted, FIRST BLOCK of packet, T011 TB means: RT PACKET, inserted, followed by a NRT BLOCK, T111 TB means: NRT, preempted, resumed, MIDDLE BLOCK of packet. T010 TB means: RT PACKET, inserted, followed by another RT PACKET, T011 TB means: RT inserted, followed by a NRT block, T101 TB means: NRT, resumed, LAST BLOCK of packet.

Figures 11 and 12 show the Preempt/Resume Transmit and Preempt/Resume Receive flow diagrams, respectively.

In these flow diagrams, a block always references a part of a NRT packet. A NRT block size is determined at the transmission side at the time the currently transmitted NRT packet is preempted.

A preempted NRT packet is reconstructed at the receive side by concatenating NRT blocks of

the same packet, received consecutively, starting with the first block, continuing with the middle block(s) if any, and terminating with the last block. First, middle and last block indications are derived from the state control bits of the Trailer Byte of each block.

A Packet references

1. A single RT packet. This type of packet is never segmented
2. A single NRT packet i.e. a non preempted packet.

In the case of a receiving sequence error, all blocks of the same NRT packet, accumulated since the first one has been received, AND the latest received block/packet are discarded.

A valid Trailer byte means the received block/packet is apparently correct:

- good X'7E' flag alignment, - checksum field of the Trailer byte validates the whole byte.

In this case the information provided by the Trailer byte is considered as good. Thanks to this information the type of the next block/packet to be received may be predicted.

An invalid Trailer byte causes the received block/packet being discarded. The type of the next block/packet to be received cannot be predicted.

The different states characterizing the system at the receive side are shown on Figure 12:

IDLE state.

This is the default state, entered at initialization time and when receiving no data from the line. IDLE state is left when receiving a first valid preempted NRT data block.

PR-RT state.

This state means the preempt-resume process has been started or is in progress. PR-RT state is entered at the end of the reception of a valid first NRT block at IDLE state or at the end of the reception of a valid middle NRT block at PR_NRT state. A RT packet is expected. The Trailer byte of any valid received RT packet indicates the following expected data is either another RT packet or a resumed NRT block. In the case of a resumed NRT block, PR_NRT state is entered.

PR_NRT state.

This state means the preempt-resume process is in progress. PR_NRT state is entered at the end of the reception of a valid last RT packet at PR-RT state. NRT block is then expected. Two cases can occur: - A middle NRT block is received, i.e. next

expected data is a RT packet then the PR-RT state is entered. - A last NRT block is received, i.e. terminating the NRT packet accumulation then IDLE state is entered.

RT state.

This state means the preempt-resume should be in progress. PR-RT state has been entered from any other state where an unexpected block/packet has been received with a valid trailer byte. This state allows the recovery of valid RT packets from a corrupted preempt/resume sequence (ex altered X'7E' end flag of a block/packet). RT state is left and IDLE state is entered as soon as a last RT packet or any non valid block/packet is received.

Figure 13 shows a block diagram of the transmitter 40 portion of the trunk interface of that communications system. Packets arrive from the communications system's packet source 41 for transmission on communication link 48. The packets may have been generated locally by this system or may have been received from another trunk on this system (e.g. an intermediate node in a packet net work). The communications system places the high-priority packets into a high priority buffer 42 and places the low-priority packets into a low priority buffer 43. If no packets are stored in either high priority buffer 42 or low priority buffer 43, a flag generator 46 is connected to the communication link 48 via a bit multiplexer 47. The flag generator 46 repeatedly generates an idle flag X'7E', when no packets are stored for transmission.

When a low-priority packet arrives in the low priority buffer 43 and a packet is currently being transmitted, the transmitter 40 waits until all earlier packets in the low priority buffer 43 have been transmitted and the high priority buffer 42 is empty. When a low-priority packet is at the head of the low priority buffer 43 and no other packet is being transmitted on the communication link 48, bytes from the low priority buffer 43 are transferred one at a time through a byte multiplexer 44 to a parallel serial converter 45. The parallel serial converter 45 serializes the data and monitors the outgoing data for sequences of five consecutive '1' bits. It also inserts a single '0' bit immediately after each set of five '1' bits. The resulting bit stream is routed through the bit mux to the communication link 48. When the transmission of the low-priority packet is complete, the bit multiplexer 47 selects the flag generator 46 for transmission to send at least one or more normal flags until the next packet is ready to be transmitted. Note that each time a flag is sent, the parallel serial converter 45 resets its internal count of the number of consecutive '1' bits.

If a low-priority packet is being transmitted from low priority buffer 43 and a high-priority pack-

et arrives in the high priority buffer 42, then the transmission of the low-priority packet is preempted. The remaining bits in the parallel serial converter 45 along with any stuffed zero bits are transmitted guaranteeing a block boundary for the preempted packet. Then the flag generator 46 sends a trailer byte, as described earlier and an idle flag X'7E'. Bytes from the high priority buffer 42 are then transferred through the byte multiplexer 44 to the parallel serial converter 45 which performs serialization and zero bit stuffing. The resulting high-priority packet is then transferred to the communication link 48. If, during the transmission of the high-priority packet, another high-priority packet arrives in the high-priority buffer 42, then the flag generator 46 sends a trailer byte and an idle flag when the first high-priority packet is completed and transmitter 40 begins transmission of the next high-priority packets without exiting the preempt mode. When the last of the series of high-priority packets has been sent (there are not more packets waiting in the high priority buffer 42), the flag generator 46 sends again a trailer byte and an idle flag. The remaining bytes from the preempted low-priority packet in the low priority buffer 43 are released to the parallel serial converter 45 and the communication link 48. If a subsequent high-priority packet arrives at the high-priority buffer 42 prior to the completion of the preempted low-priority packet, the preemption and resume sequence is repeated. When the transmission of the low-priority packet is completed, the flag generator 46 transmits an idle flag.

Figure 14 shows a block diagram of the receiver 50 portion of the communication link interface of the communications system up to a point at which received whole packets are passed to a packet target 56 within the communications system. The packet target 56 could be the final destination for the received packets or could be a packet switch used to route packets to other trunks for transmission to other nodes in a packet network. Any buffering associated with the packet target 56 is outside the receiver 50 and is not included in Figure 14.

A flag detector 52 continuously monitors the bit stream received from a communication link 51 trailer byte and idle flags. If a sequence of bits other than a flag is detected immediately following an idle flag, it indicates the beginning of a new frame. A serial parallel converter 53 receives the bit stream, discard any '0' bit if it immediately follows five consecutive '1' bits, and converts the remaining bits into byte-parallel form. If the received packet is a high-priority packet, the parallel byte data is passed directly through a multiplexer 59 to a multiplexer 55 connected to the packet target 56 until a normal ending flag is detected by the flag

detector 52. The receiver 50 indicates the end of the packet to the packet target 56.

If the received packet can be preempted (i.e. low-priority, non-real-time packet), then the parallel byte data is instead passed through byte multiplexer 59 to the preemptable packet buffer 54 in order to permit the entire packet to be accumulated before passing it to the packet target 56. If flag detector 52 detects a trailer byte that there is no partial byte in the serial parallel converter 53, then it indicates the beginning of a high-priority preempting packet and therefore the beginning of preempt mode. The bit stream is passed through the serial parallel converter 53 as before but this time the parallel byte data is passed directly through the multiplexers 59 and 55 to the packet target 56 within the communications system.

When the flag detector 52 detects either a normal ending flag or a trailer byte indicating the end of preemption, then the receiver 50 indicates the end of the packet to the packet target 56. If a normal ending flag is detected, then the serial parallel converter 53 will continue to route the parallel byte data from subsequent packets directly to the multiplexer 55. If a trailer byte indicating the end of preemption is detected, the receiver 50 will end preempt mode. The received bit stream will be routed through the serial parallel converter 53 and multiplexer 59 to the preemptable packet buffer 54 thus resuming reception of the preempted low-priority packet. If the flag detector 52 detects a normal ending flag indicating the end of the low-priority packet, the receiver 50 transfers the entire low-priority packet stored in the preemptable packet buffer 54 through the multiplexer 55 to the packet target 56.

Claims

1. A packet-switched communications system for transmitting low-priority packets segmented in short blocks and high priority packets characterized in that it comprises :
 - means for embedding high-priority packets within low-priority packets based on distinctive bit patterns attached to each block of said low-priority packets and to each of said high-priority packets;
2. A system according to claim 1 wherein said means for embedding high-priority packets within low-priority packets comprises :
 - means responsive to a first distinctive bit pattern for transmitting the high-priority packet before the next block of low-priority packet,
 - means for buffering the next blocks of the low-priority packet that is preempted

- by the high-priority packet, and
- means responsive to a second distinctive bit pattern for resuming transmission of the blocks of the low-priority packet temporarily stored in said buffering means.

5

3. A system according to claim 2 further comprising means for checking whether or not a high-priority packet is to be transmitted.

10

4. A system according to claim 3 further comprising means for prefixing each packet with a starting flag.

5. A system according to anyone of claims 1-4 wherein said distinctive bit pattern is a trailer byte.

15

6. A method for embedding high-priority packets within a low-priority packet segmented in short blocks in a packet-switched communications system, characterized in that it comprises the steps of :

20

- a. prefixing each low-priority packet with a starting flag ;
- b. suffixing each low-priority packet with a distinctive bit pattern ;
- c. suffixing each high-priority packet with a distinctive pattern of bits ;
- d. using the starting flag to initiate transmission of packets and to establish byte alignment with a packet frame ;
- e. using a first distinctive pattern of bits to interrupt the transmission of low-priority packets and buffer untransmitted low-priority while transmitting a high-priority packet ;
- f. using a second distinctive pattern of bits following the completion of transmission of a high-priority packet to indicate the resumption of the transmission of the interrupted low-priority packet ; and
- g. using the ending flag to return the communications system to an idle state.

25

30

35

40

45

50

55

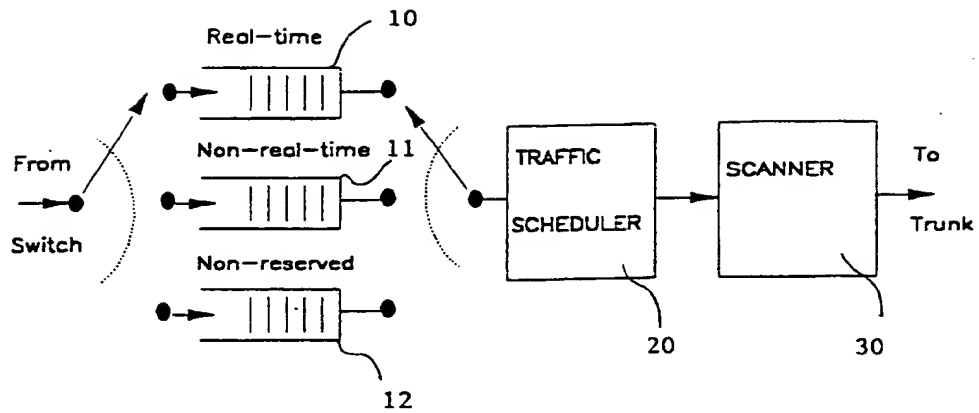


FIG. 1

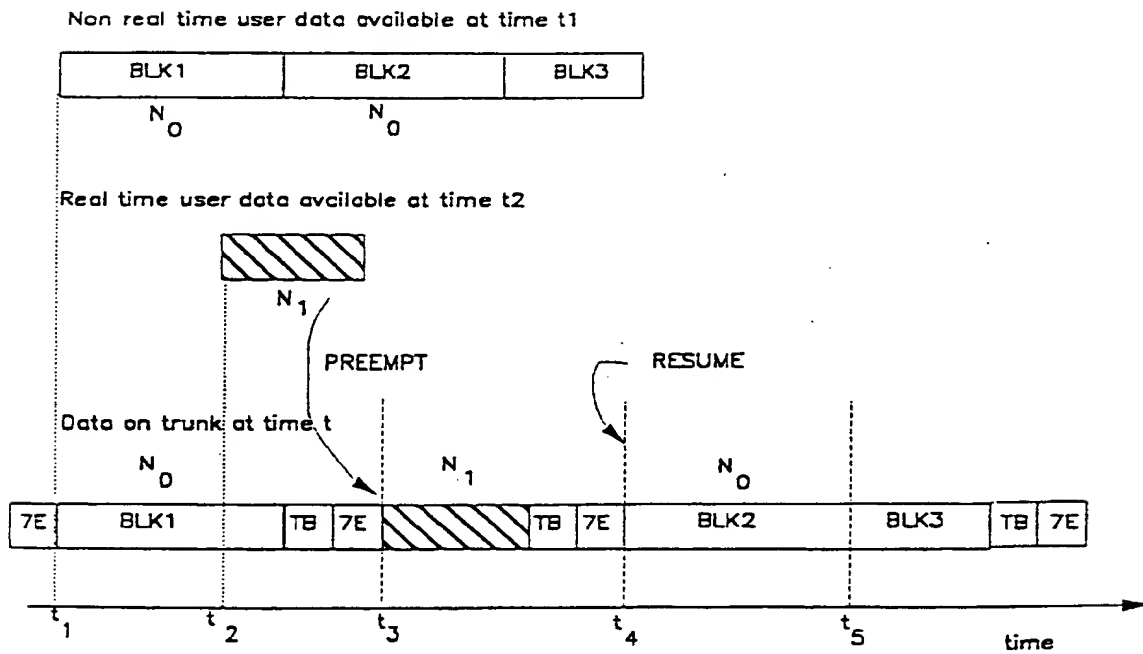


FIG. 2




FIG. 3




FIG. 4

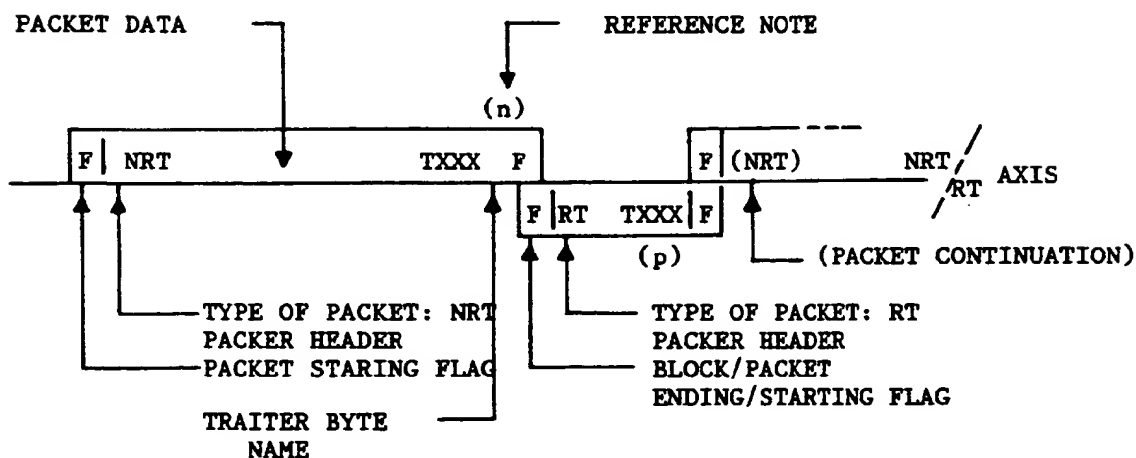


FIG. 5




FIG. 6




FIG. 7




FIG. 8




FIG. 9




FIG. 10

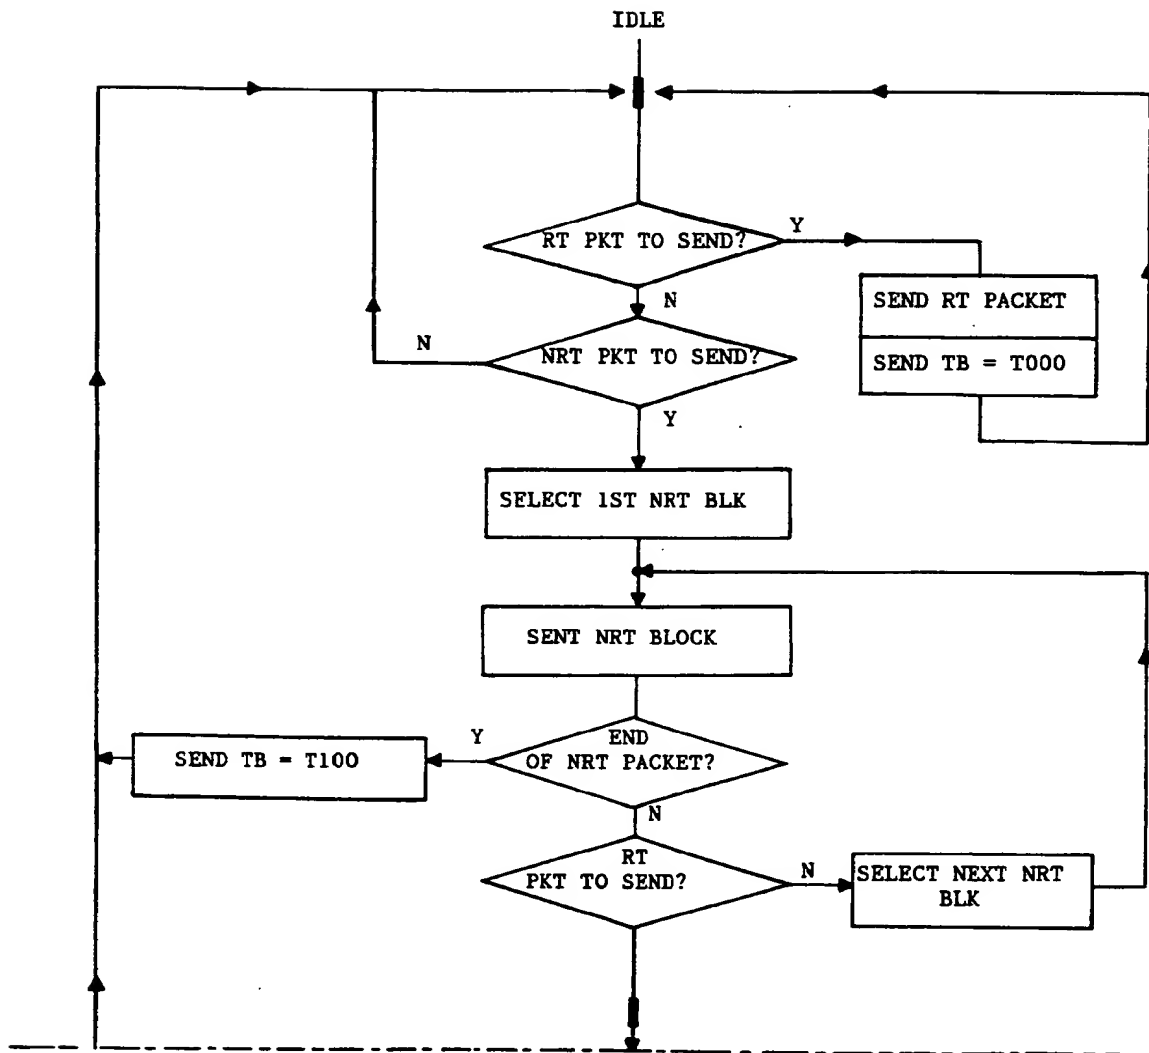


FIG. 11A

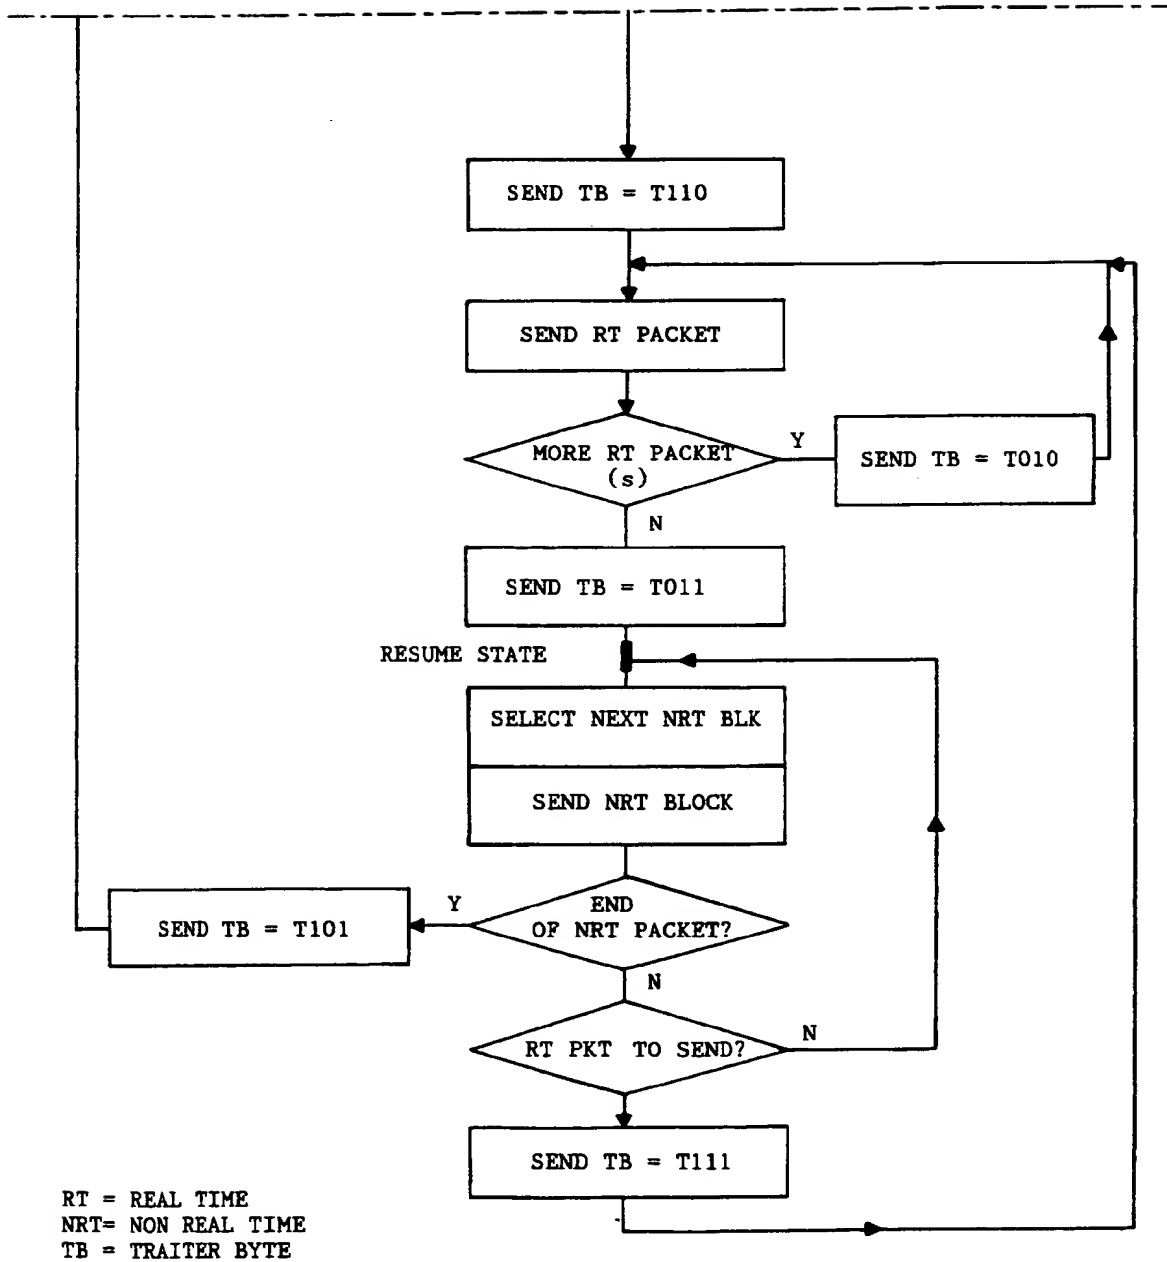


FIG. 11B

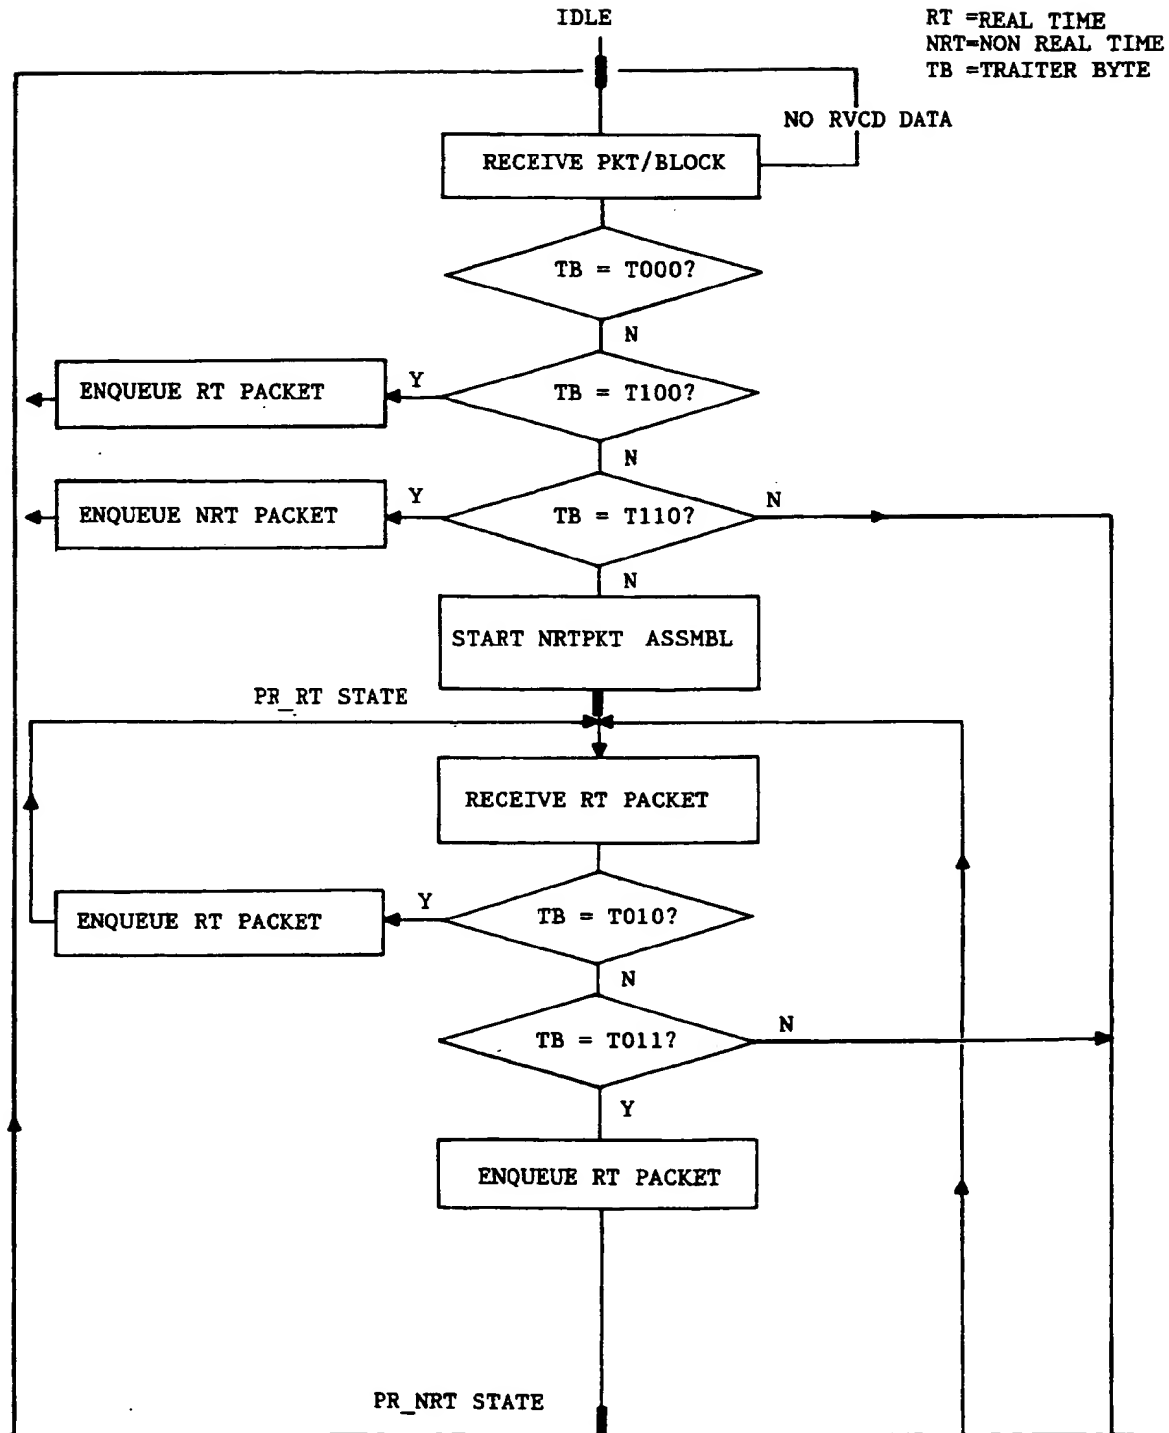


FIG. 12A

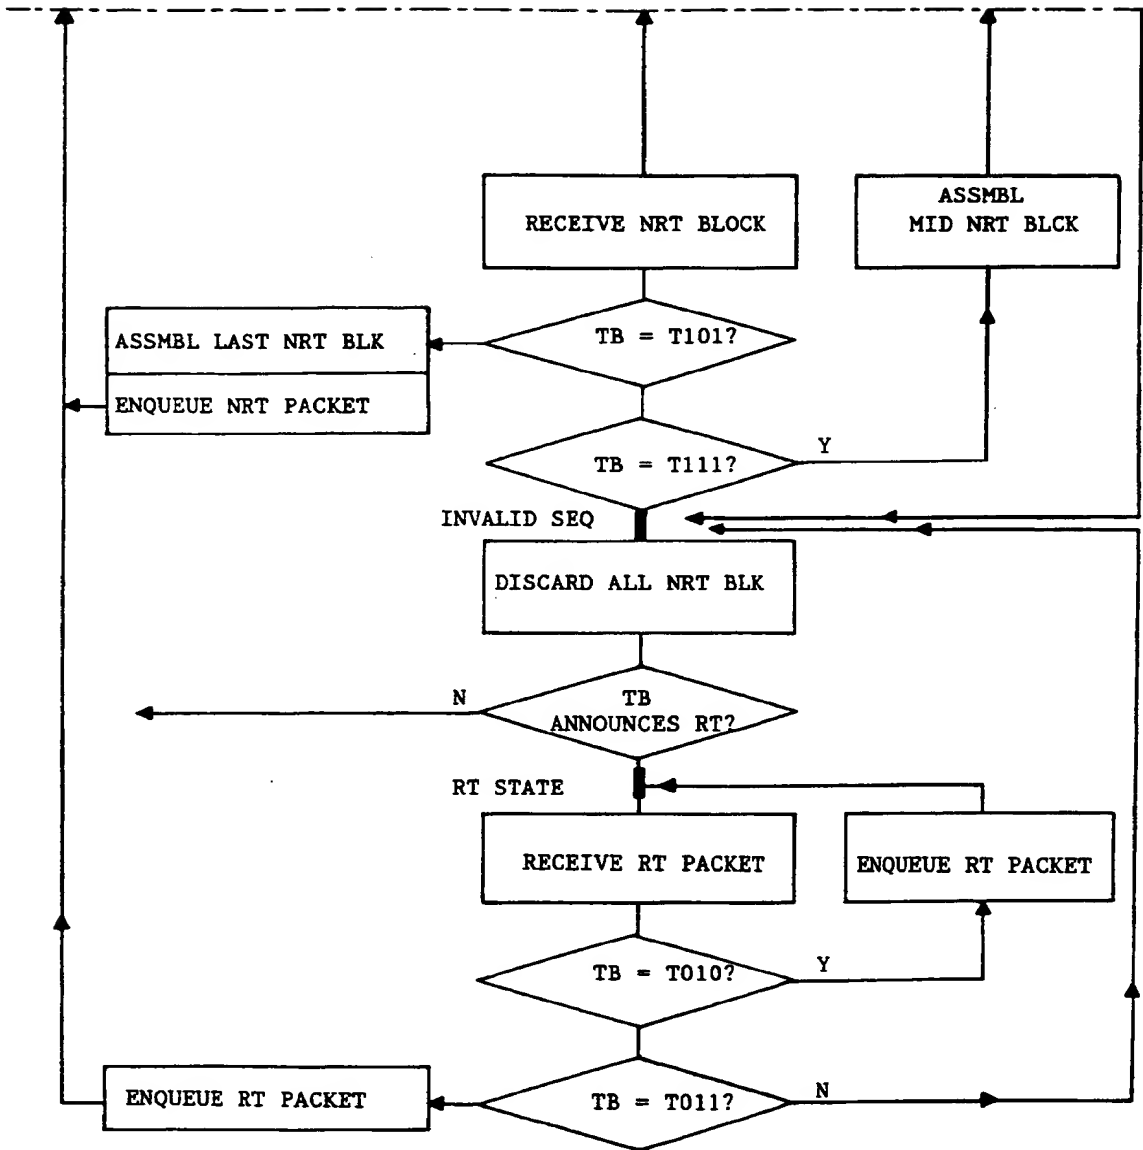




FIG. 12 B





European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 94 48 0047

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
D,X	EP-A-0 582 537 (INTERNATIONAL BUSINESS MACHINES) * page 2, line 39 - line 48 * * claims 1-8 * * figures 2,3 * ---	1-6	H04L29/06
A	US-A-4 510 599 (E.U. MEHMET) * abstract * * claims 21-23,28-30,32 * ---	1-6	
A	IBM TECHNICAL DISCLOSURE BULLETIN, vol.37, no.2A, February 1994, NEW YORK US pages 243 - 245 'MAC layer handling of preemptive and nonpreemptive priorities in buffer insertion LANs' * the whole document * ---	1-6	
A	EP-A-0 512 290 (INTERNATIONAL BUSINESS MACHINES) * abstract * * page 7, line 45 - line 58 * -----	1-6	
			TECHNICAL FIELDS SEARCHED (Int.Cl.6)
			H04L
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 2 December 1994	Examiner Perez Perez, J
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document			

THIS PAGE BLANK (USPTO)